# Final Report - Extending Cooperative SLAM into Multi-Objective Missions

Mitch Bryson and Salah Sukkarieh

## I. Introduction

This final report discusses the theoretical and practical outcomes of the second half of the project "Extending Cooperative SLAM into Multi-Objective Missions". This project aims to extend the work of cooperative SLAM in [1], [6] into missions involving multiple objectives such as directing a group of aerial vehicles to a specific destination in minimum time. The strong interdependencies between the tasks of control, localisation, mapping and exploration create tight coupling between multiple objectives in the mission. Additionally, because prior information about the terrain is partial, the planning and control tasks are required to be robust to variations in the environment with the ability to adapt to new information as it becomes available. This project aims to study these interdependencies by formulating guidance and path planning algorithms for UAVs tasked to reach a goal destination over partially-known terrain.

In the first half of this project [2], we developed a framework for both single and multi-UAV path planning over Global Positioning System (GPS) denied, partially-known terrain. The framework was based on an $A^*$ search algorithm [4] which found a sub-optimal minimum-time path between two locations over the terrain while ensuring that several constraints based on the localisation accuracy in the UAV's terrain-aided navigation system were met. Simulation results of the path planning scheme showed that the method was able to plan feasible paths that met both the localisation error and time constraints and balanced between traversing known and unknown areas of the terrain in a measured, quantitative way, that was dependant on the magnitude of the constraints and the expected density of features. In this second half of the report, we focus on improving key areas in the path planning system that are applicable to the single and multi-vehicle versions of the algorithms developed in [2]. Of the several avenues for continuing work, two key areas were identified from the results:

1) Estimating terrain-aided localisation system performance while flying over unmapped areas of the terrain required an approximation of the density of terrain features in this area and thus resulted in a degree of uncertainty of the performance. It was proposed that the magnitude of this uncertainty and it's repercussions to the planning process should be studied, along with methods for accounting for the uncertainty or risk in paths planned over unmapped terrain.
2) In the first half of the project, the proposed planning methods were limited to initial path planning at the start of a UAV mission. It was proposed that the methods should be extended to provide the ability to replan and alter paths on-line, while the vehicle is in-flight based on new terrain information that becomes available, providing a robust and adaptable planner.

In this report we discuss the theoretical and practical results of these research questions and the work completed in the second half of the project. In Section II we examine the first research question by providing an analysis of the variance in the expected localisation performance of the path planning scheme and examine the effect of traversing known/unknown terrain and terrain feature density. We also propose a path planning scheme for accounting for uncertainty in path performance. In Section III we examine the second research question by extending the path planning architecture in [2] to include the ability to replan the vehicle's path based on new information received while flying the trajectory. Results are presented in different scenarios where the actual density of features discovered in the terrain varies from estimates made at the beginning of the path planning process. In both these sections, results are presented from a 6-DoF simulation of a UAV flying over the terrain, with sensor measurements simulated from three

# Report Documentation Page

*Form Approved*
*OMB No. 0704-0188*



| 1. REPORT DATE | 2. REPORT TYPE | 3. DATES COVERED |
|---|---|---|
| **05 MAY 2009** | **FInal** | **18-04-2008 to 17-09-2008** |

4. TITLE AND SUBTITLE
**Extending Cooperative SLAM for Multi-Objective Missions**

5a. CONTRACT NUMBER
**FA48690814060**

5b. GRANT NUMBER

5c. PROGRAM ELEMENT NUMBER

6. AUTHOR(S)
**Sukkarieh Salah**

5d. PROJECT NUMBER

5e. TASK NUMBER

5f. WORK UNIT NUMBER

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)
**University of Sydney,University of Sydney, Bldg. J04,Sydney NSW 2006,Australia,au,2006**

8. PERFORMING ORGANIZATION REPORT NUMBER
**N/A**

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)
**AOARD, UNIT 45002, APO, AP, 96337-5002**

10. SPONSOR/MONITOR'S ACRONYM(S)
**AOARD**

11. SPONSOR/MONITOR'S REPORT NUMBER(S)
**AOARD-084060**

12. DISTRIBUTION/AVAILABILITY STATEMENT
**Approved for public release; distribution unlimited**

13. SUPPLEMENTARY NOTES

14. ABSTRACT
**This report discusses theoretical and practical outcomes of the second half of a project that extended the work of cooperative SLAM into missions involving multiple objects, such as directing a group of aerial vehicles to a specific destination in minimum time**

15. SUBJECT TERMS
**Cooperative Control, Decentralized Sensor Networks, Simultaneous Localization And Mapping**

| 16. SECURITY CLASSIFICATION OF: | | | 17. LIMITATION OF ABSTRACT | 18. NUMBER OF PAGES | 19a. NAME OF RESPONSIBLE PERSON |
|---|---|---|---|---|---|
| a. REPORT | b. ABSTRACT | c. THIS PAGE | | | |
| **unclassified** | **unclassified** | **unclassified** | **Same as Report (SAR)** | **21** | |

**Standard Form 298 (Rev. 8-98)**
Prescribed by ANSI Std Z39-18

on-board terrain sensors and an on-board Inertial Measuring Unit (IMU). In Section IV we summarise and discuss the conclusions of the research and suggest areas for future work in Cooperative SLAM.

## II. Accounting for Uncertainty in Localisation Performance over Unmapped Terrain

In this section we begin by examining the variation in expected localisation accuracy when flying over unknown terrain based on the method proposed in [2]. From this examination, we propose an addition to the path planning algorithms in [2] which account for this uncertainty.

### A. *Overview of SLAM Localisation Accuracy Prediction*

In order to make localisation accuracy-optimising decisions about prospective paths a UAV could traverse over partially known terrain, a path planner must have a method for predicting the expected localisation performance associated with flying a given path. When flying over the known sections of the terrain, an estimate of the positions of map features is known by the planner beforehand, and thus, based on the known configuration of landmark sensors on-board the UAV, the path planner can predict the observations that will be made by the sensors and thus predict the terrain-aided localisation system performance from these expected measurements. When traversing the unknown sections of the map, the UAV will find new features (the positions of which are unknown before flying) in the terrain and use these features in a Simultaneous Localisation And Mapping (SLAM) estimation framework to estimate the UAV's self-location along the trajectory. Unfortunately the locations of these features is not known before flying; instead all that is known is an expected density of available features based on the other known sections of the map, which complicates the process of predicting the localisation system performance associated with these paths.

In [2], a method was proposed for predicting localisation performance over unknown sections of terrain by randomly distributing 'expected' feature locations into the unknown sections of the map based on the density of features present in the known sections of the map. These 'expected' feature locations were then used to produce expected landmark sensor observations for a given trajectory which were used to predict the SLAM localisation accuracy. When the actual path is flown by the vehicle, the actual configuration of features observed would be different, however the density of features would be the same and thus it was hoped that the predicted SLAM localisation accuracy would be a good approximation to the actual localisation accuracy.

### B. *Performance Prediction Variance Results*

In order to test the accuracy of SLAM localisation performance prediction method described above, we computed the variance in predicted performance along a set vehicle trajectory when the random distribution of expected features was varied. A simple straight line trajectory between two sets of prior-mapped features was chosen to examine the variation in performance prediction vs. terrain feature density. Figure 1 illustrates the 2km long trajectory and positions of prior-mapped features.

Several instances of random feature maps were then generated, where the expected features were distributed into the unknown areas of the map. For each random map, the expected localisation accuracy of the trajectory over the terrain was predicted based on the method presented in [2]. Four accuracy metrics (yaw angle, roll/pitch angle, horizontal positioning and vertical positioning accuracy) based on the expected SLAM localisation covariance were predicted along the trajectory. Experiments were conducted at different expected terrain feature densities (from 2 to 10 features per 100x100m), and for each terrain density, the accuracy metrics were predicted along the trajectory over 30 different random permutations of the expected map locations. For each terrain density, the average, variance and maximum variation of the predicted accuracy metrics along the trajectory for the 30 different random maps were recorded.

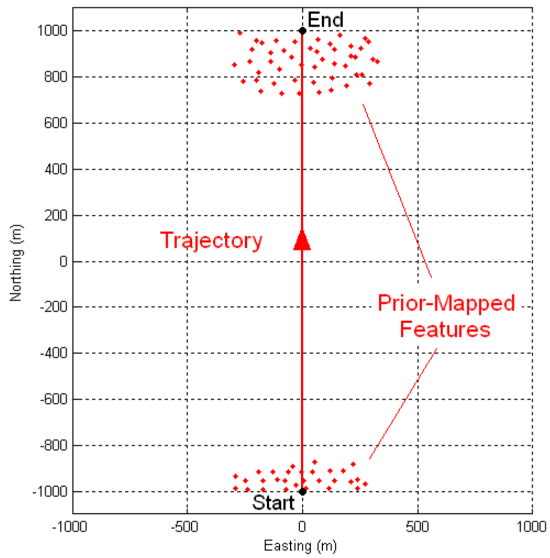


Fig. 1. Simple Straight and Level Trajectory for Examining Path Performance Prediction Variations: The path consists of a single 2km long path of straight and level flight between two groups of prior known features shown in red. While flying over the unmapped area, the vehicle will use SLAM to help constrain localisation errors. The predicted localisation performance is evaluated for the trajectory for several different values of assumed terrain feature density in the unmapped area.

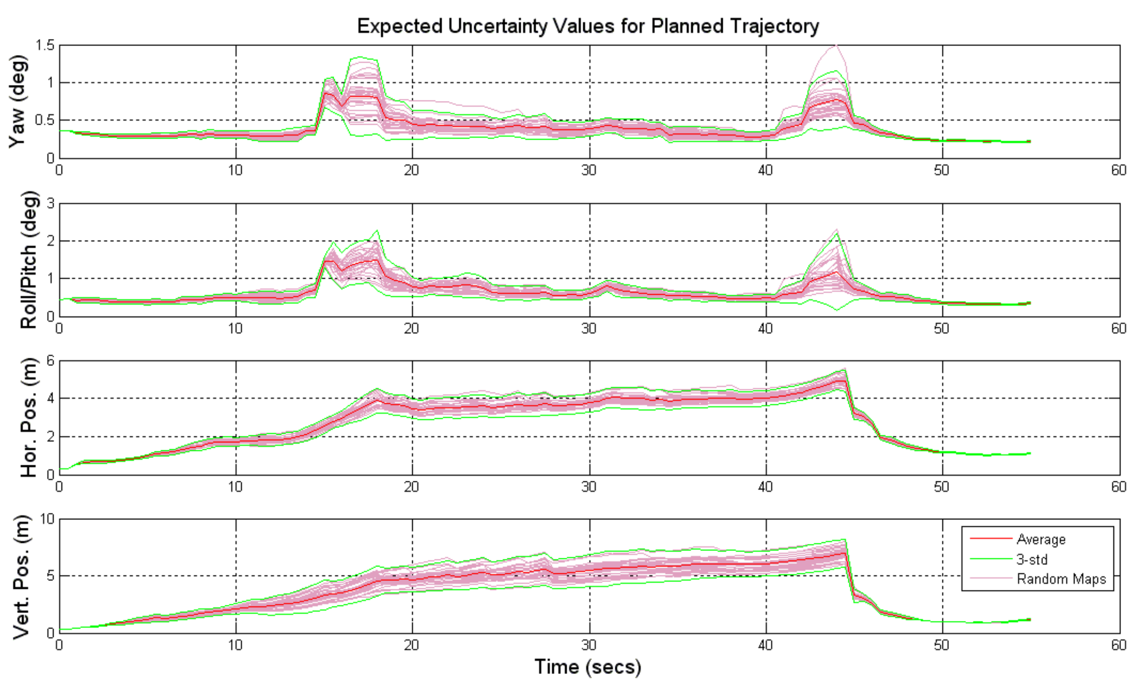


Fig. 2. Variation in Predicted Localisation Accuracy Metrics for Low Density Expected Terrain Features: Shown in red is the predicted value of each of the four performance metrics at each point along the trajectory, averaged out over 30 simulation runs. The green bounds indicate the $3\sigma$ variation of the predicted performance and the magenta lines indicate the performance predicted by each of the 30 runs. A large variation in the predicted performance is present in all the metrics, mostly during the middle of the trajectory, where the prediction applies to the vehicle operating over unknown terrain.

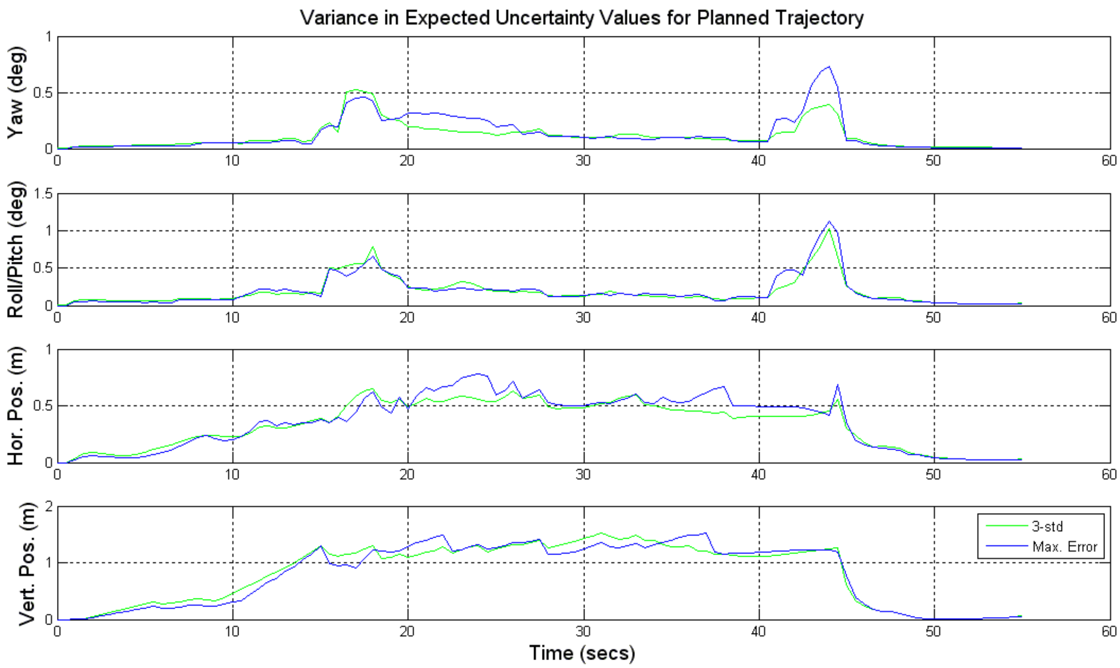


Fig. 3. Variation in Predicted Localisation Accuracy Metrics for Low Density Expected Terrain Features: Shown in green is the $3\sigma$ variance of the each of the four predicted localistion accuracy metrics and shown in blue is the maximum variation from the mean out of the 30 simulation runs performed. The maximum variation is reasonably close the $3\sigma$ variance, except perhaps for the yaw and roll/pitch metrics (for example near the 43 second mark) where the maximum variation is quite high (see discussion in text).

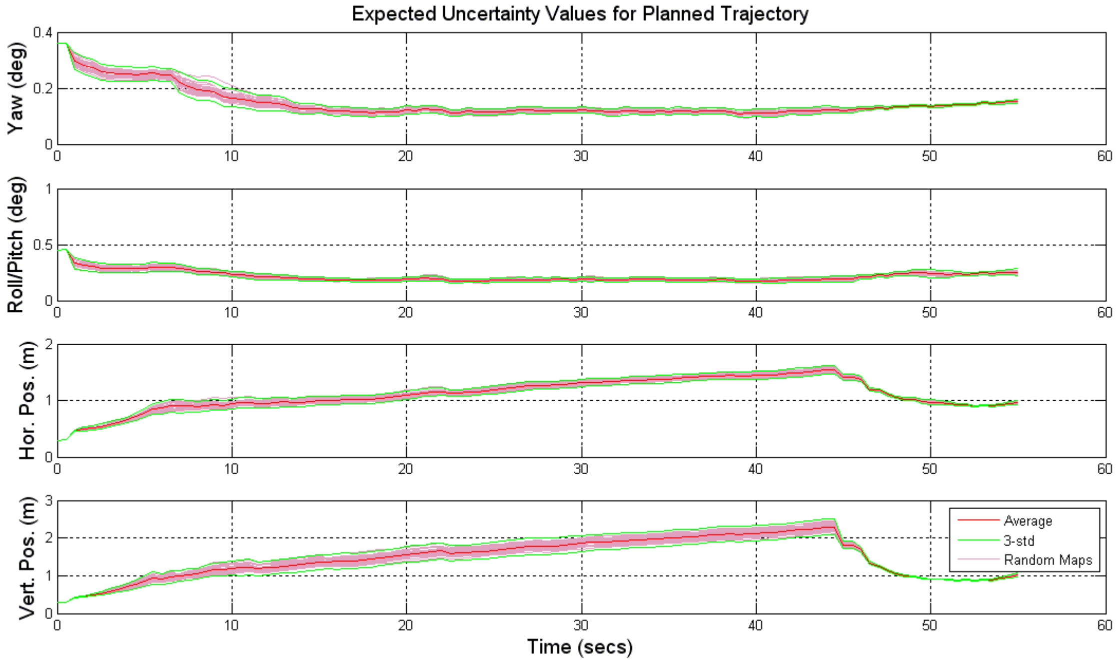


Fig. 4. Variation in Predicted Localisation Accuracy Metrics for High Density Expected Terrain Features: Shown in red is the predicted value of each of the four performance metrics at each point along the trajectory, averaged out over 30 simulation runs. The green bounds indicate the $3\sigma$ variation of the predicted performance and the magenta lines indicate the performance predicted by each of the 30 runs. The variation in the predicted metrics is now much less than the low density feature case.

Figures 2 and 3 illustrate the variation in the four performance metric predictions along the trajectory over the 30 simulation runs where the expected feature density is low (2 features per 100x100m area). Shown in the plots is the mean performance metric value, $3\sigma$ variance and each of the 30 predicted values. It can be seen from the plots that the variance in the prediction performance is reasonably small during the beginning and end of the trajectory (during the traversal of prior-known terrain features) and higher during the middle section of the trajectory (where the vehicle traverses unmapped sections of the terrain). The low density of terrain features means that only a small number of features may be present in the UAV's landmark sensor field of view at any instance in time. The exact SLAM algorithm performance is highly dependant on the position of the features w.r.t the UAV and for times where only 1 or 2 features are present in the field of view, the predicted SLAM performance becomes sensitive to exactly where this feature is observed.

Figures 4 and 5 illustrate the variation in the four performance metric predictions along the trajectory over the 30 simulation runs where the expected feature density is higher (10 features per 100x100m area). The variation in performance prediction is now much lower than the 2 features per 100x100m area case. At any instance in time while operating over unmapped terrain, the UAV terrain sensor can expect to see in the order of 10 features. Since these features tend to be spread evenly across the sensor field of view, the SLAM predicted performance is much less sensitive to the exact positions of the features.

Figure 6 illustrates the variations in performance metrics, averaged out over sections of the trajectory where the vehicle traverses known features (start and end of trajectory) and averaged out over sections of the trajectory where the vehicle traverses unmapped terrain (middle of the trajectory). The values are compared for several different expected terrain densities. In general, the performance prediction variance is lower when operating over known sections of the map, than unknown sections of the map, as the positions of these features are reasonably well-known beforehand. When operating over unknown sections of the terrain, the performance prediction variance decreases asymptotically as the expected terrain density increases. It is also noted that as feature density is decreased, performance prediction variance rises slightly

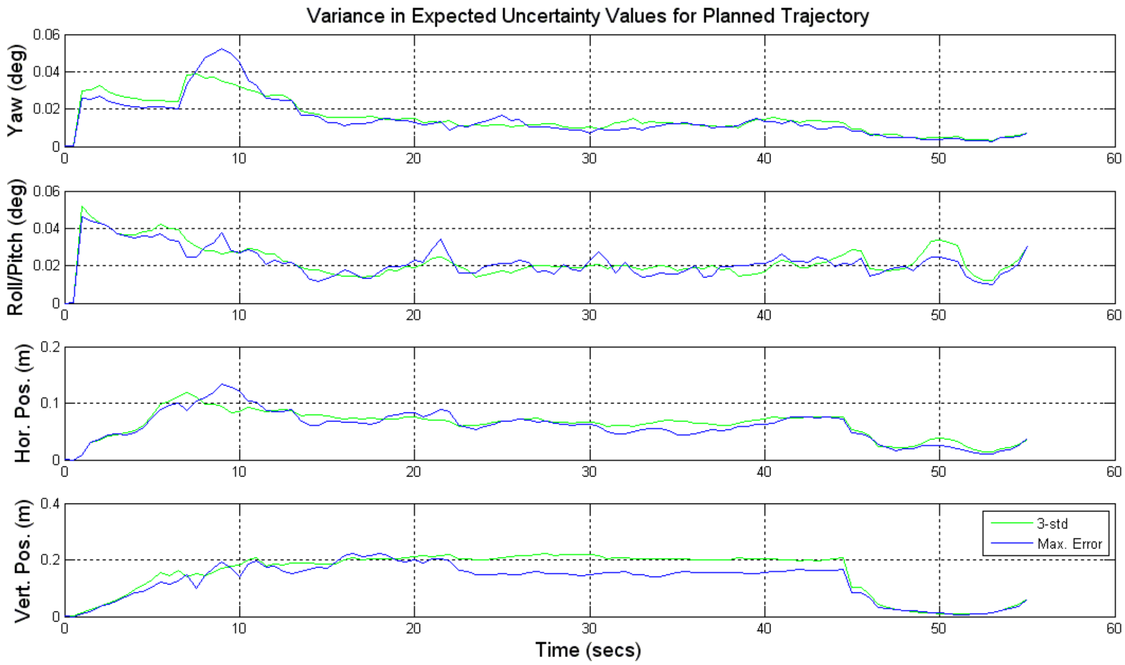


Fig. 5. Variation in Predicted Localisation Accuracy Metrics for High Density Expected Terrain Features: Shown in green is the $3\sigma$ variance of the each of the four predicted localistion accuracy metrics and shown in blue is the maximum variation from the mean out of the 30 simulation runs performed. The variation in the predicted metrics is now much less than the low density feature case.

more sharply in the angular-based performance metrics (yaw and roll/pitch) vs. the position-based metrics. This is due to the fact that there are small sections of the flight where the feature density is low enough such that there are no features observed at a particular instance in time. Although the vehicle attitude estimate uncertainty does not grow as quickly as the position estimate uncertainty does during un-aided inertial flight, the variation in attitude uncertainty is larger at these times owing to relative orientation of feature observations w.r.t the UAV playing a larger role in SLAM estimate observability (see [6], [3] for further details).

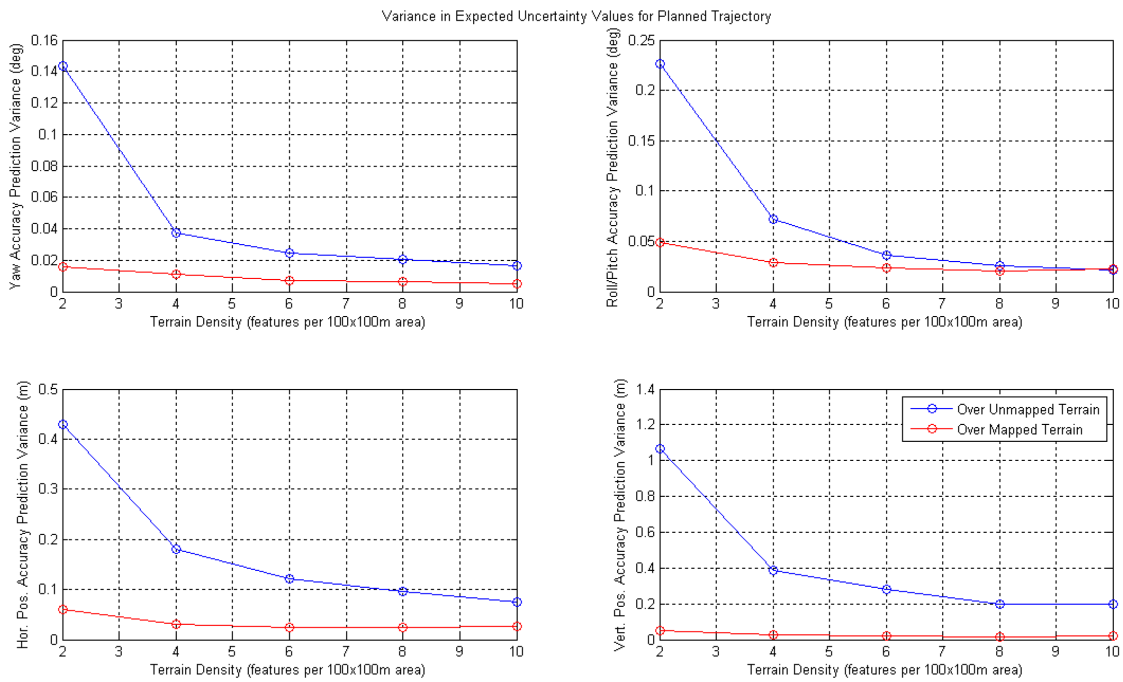


Fig. 6. Comparison of Predicted Localisation Accuracy Metrics over Different Values of Expected Terrain Density: The blue line indicates the performance metric variance averaged over trajectory times where the vehicle traverses unknown terrain, for different expected terrain density values. The red curve indicates the performance metric variance averaged over trajectory times where the vehicle traverses known terrain features. The variation in performance prediction is lower when operating over known terrain, and when operating over unknown terrain, the variation is higher when the expected feature density is low.

### C. Accounting for Performance Uncertainty over Unmapped Terrain

The results of the previous subsection demonstrate the effectiveness of the path planning method proposed in the first half of the project [2] when operating over unknown terrain where the expected terrain density is high (relative to the terrain sensor footprint of each UAV). In this case, the variation in performance prediction is low and thus a good approximation of the localisation performance can be made from a single random instance of the expected map. However, when the expected terrain density is low, there is a higher degree of variation in the predicted performance. This could result in the path planning scheme 'over-estimating' the expected localisation accuracy associated with a planned path, such that when the actual path is flown the localisation accuracy is lower and the original accuracy constraints applied to the planner are violated. In order to avoid violating path localisation performance constraints, the path planner should account for the inherent risk or uncertainty associated with flight over low-density unmapped terrain.

*1) Constraint Variation for Uncertain Path Planning:* The $A^*$ search-based planning algorithm presented in [2] assesses the feasibility of potential plans by assessing whether the expected localisation errors will rise above specified thresholds during any part of the flight. Four constraint values are considered: yaw angle ($e_{yaw} < c_{yaw}$), roll/pitch angle ($e_{roll,pitch} < c_{roll,pitch}$), horizontal positioning ($e_{hor.pos} < c_{hor.pos}$) and vertical positioning ($e_{vert.pos} < c_{vert.pos}$) errors where $e_{yaw}$, $e_{roll,pitch}$, $e_{hor.pos}$, $e_{vert.pos}$ are the expected errors ($3\sigma$ expected SLAM covariance) and $c_{yaw}, c_{roll,pitch}, c_{hor.pos}, c_{vert.pos}$ are the constraint thresholds. The constraint thresholds are all set at fixed values $c_{0,yaw}$, $c_{0,roll,pitch}$, $c_{0,hor.pos}$, $c_{0,vert.pos}$ at the start of the planning (dictated by external objectives of the UAV's mission). In the remaining results in this paper it is assumed that the constraint variables are set at $c_{0,yaw} = 3^o/s$, $c_{0,roll,pitch} = 3^o/s$, $c_{0,hor.pos} = 10m$ and $c_{0,vert.pos} = 5m$. The proposed risk-adverse planning system varies the constraints to $c_{i,yaw}, c_{i,roll,pitch}, c_{i,hor.pos}, c_{i,vert.pos}$ for a given section $i$ of the trajectory:

$$
\begin{aligned}
c_{i,yaw} &= c_{0,yaw}(1 - k.r_{obs}) && (1)\\
c_{i,roll,pitch} &= c_{0,roll,pitch}(1 - k.r_{obs}) && (2)\\
c_{i,hor.pos} &= c_{0,hor.pos}(1 - k.r_{obs}) && (3)\\
c_{i,vert.pos} &= c_{0,vert.pos}(1 - k.r_{obs}) && (4)
\end{aligned}
$$

where $r_{obs}$ is the ratio of the number of expected unmapped feature observations $z_{unmapped}$ to total number of expected feature observations (including both unmapped and prior known features observations $z_{known}$) during the trajectory segment $i$:

$$
r_{obs} = \frac{\sum z_{unmapped}}{\sum(z_{unmapped}, z_{known})} \tag{5}
$$

and $k$ is a tuning uncertainty constant, based on the expected density of features:

$$
\begin{aligned}
k &= 1 - \tfrac{\rho}{\rho_s}, \quad \rho \leq \rho_s\\
&= 0, \qquad\quad \rho > \rho_s
\end{aligned} \tag{6}
$$

where $\rho$ is the expected density of features in the map and $\rho_s$ is the 'safe' density of features where the expected performance variance vs. density curve drops below a user-specified level.

*2) Results:* The proposed risk-adverse planning system was compared to the original path planning scheme presented in [2] over a 6km trajectory. The expected density of terrain features in unmapped areas was set at 2 features per 100x100m area and the 'safe' feature density value $\rho_s$ was set to 3 features per 100x100m area (based on the results of Figure 6). Note that both planned trajectories were made with the same random set of expected features placed into the unmapped areas.

Figures 7 (a) and (b) illustrate the paths taken by the vehicle for the both the original, non-risk adverse path planner and the extended, risk adverse planning scheme with online adjusted constraint values. As is consistent with prior results in [2], the planned trajectory performs a kind of 'looping' or doubling-back' strategy for traversing unmapped areas in order to form continuous, small loop closures in SLAM which

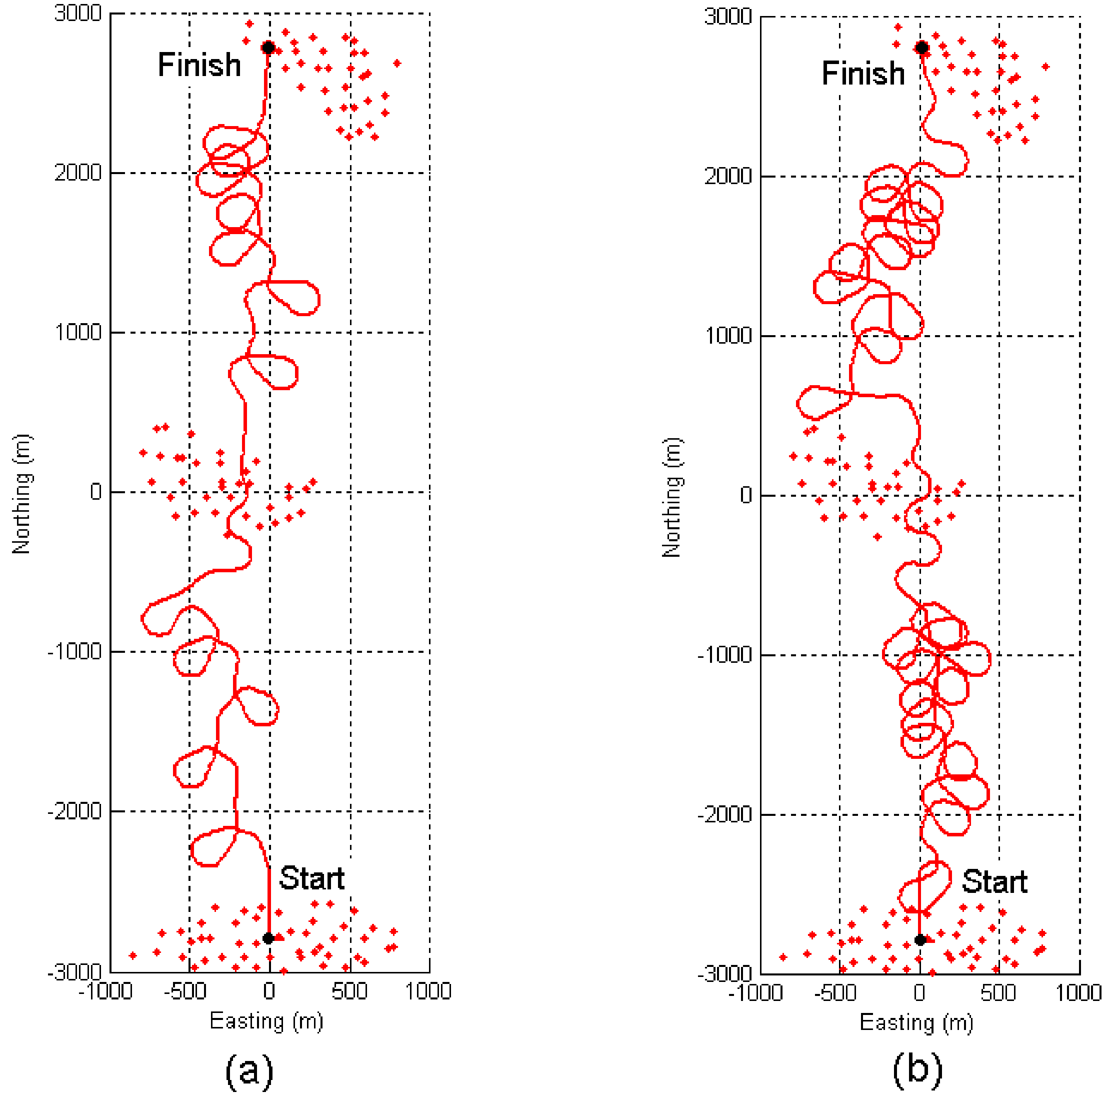


Fig. 7. Planned Vehicle Trajectories over Example Terrain Set: (a) Planned vehicle trajectory for the original, non-risk adverse planner and (b) Planned vehicle trajectory for the risk adverse planner. The vehicle trajectories are indicated by the red line and the red points indicates the position of prior known map features. The risk adverse trajectory more conservatively 'doubles-back' on the trajectory through unmapped regions while planning a similar path to the non-risk adverse planner over prior mapped regions.

constrain localisation error buildup while traversing unmapped areas. While traversing the prior mapped segments of the terrain, the flight trajectory is straight and level, heading towards the destination due the the ability to use the terrain information present.

It can be seen in the risk-adverse path (Figures 7 (b)) that the vehicle performs a larger number of 'doubling-back' maneuvers, with greater frequency over the unmapped areas. This path therefore builds up more features observations and moves into the unmapped areas more gradually than in the non-risk adverse case, a safer strategy that accounts for uncertainty involved in traversing this part of the terrain. When the vehicle trajectory passes over the middle, prior-known section of the terrain, the path is much less conservative, owing to the fact that the position of features is reasonably well known (from the prior map information) and thus the uncertainty involved in the performance prediction is less.

Figure 8 illustrates the variation in predicted localisation performance metrics over the course of the trajectory for the non-risk adverse path planner. As can be seen from the figure, the variance in performance prediction is higher for sections of the trajectory where the vehicle traverses over unmapped sections of

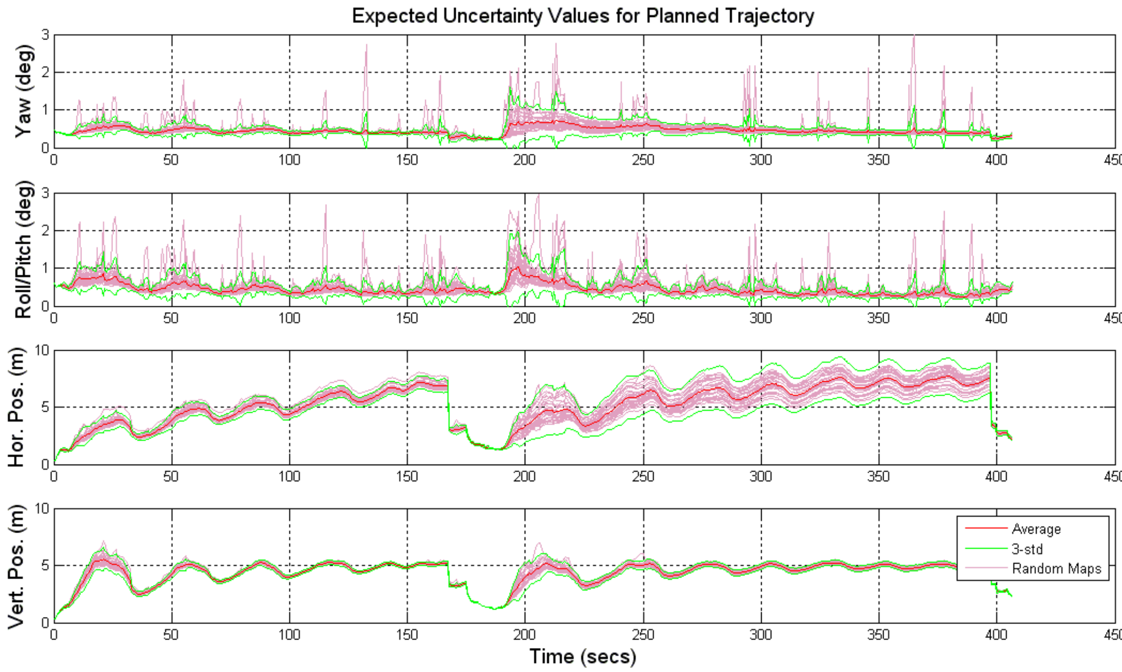


Fig. 8. Variation in Predicted Localisation Accuracy Metrics for the path taken by the original, non-risk adverse path planning system: Shown in red is the predicted value of each of the four performance metrics at each point along the trajectory, averaged out over 30 simulation runs. The green bounds indicate the $3\sigma$ variation of the predicted performance and the magenta lines indicate the performance predicted by each of the 30 runs. Although the average predicted performance metric values remain below the specified thresholds applied by the $A^*$ path planning algorithms, there exist map configurations where the predicted error levels rise above the constraints, indicating a 'risk' in traversing the path.

the terrain (i.e. 5 to 170 seconds and 190 to 395 seconds). It is also noticed that, particularly for the yaw and roll/pitch metrics, that, for a small number of the simulation runs, the performance variation rises sharply out of the $3\sigma$ bounds (purple peaks in the plot). This occurs in some random configurations of the map where there are sections of the terrain with one or no features for a short period of time (owing to the low density of expected feature locations). These sharp peaks mean that there is a certain risk that if the 'true' configuration of features in the area has similar gaps, then the constraints on the localisation performance metrics during the flight may be broken.

Figure 9 illustrates the variation in predicted localisation performance metrics over the course of the trajectory for the risk adverse path planner. The variance in the predicted values of localisation errors has now been reduced (along with the mean expected values themselves) such that the localisation error constraints are never violated in any of the 30 simulation runs of the planned trajectory. There is of course a cost associated with the risk adverse path; the total time taken to traverse the terrain has now increased from 410 seconds to 660 seconds owing to the more conservative nature of the path. It should be noted however that the path is not simply more conservative in general; the conservative nature of the path is scaled to the uncertainty in predicted performance. In situations where the vehicle moves across prior known terrain, the path taken is the same as for the non-risk adverse case, where risk of constraint violation is much smaller.

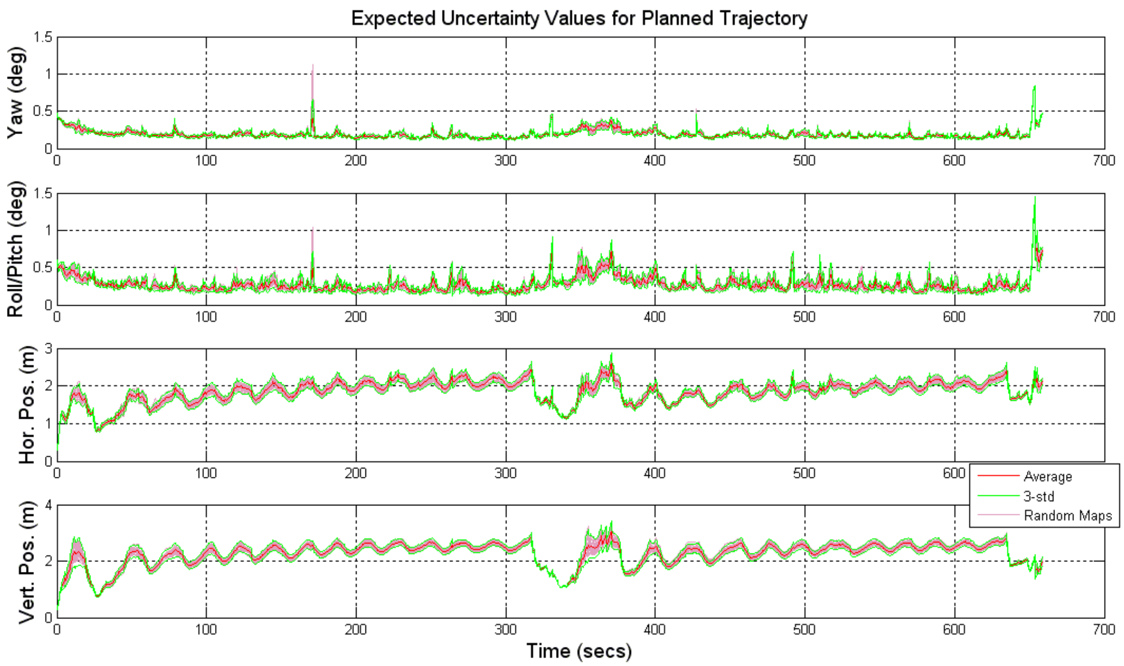


Fig. 9. Variation in Predicted Localisation Accuracy Metrics for the path taken by the risk adverse path planning system: Shown in red is the predicted value of each of the four performance metrics at each point along the trajectory, averaged out over 30 simulation runs. The green bounds indicate the $3\sigma$ variation of the predicted performance and the magenta lines indicate the performance predicted by each of the 30 runs. Note now that the localisation constraints are never violated in any of the 30 simulation runs, indicating a much safer trajectory, but at the cost that the trajectory has taken longer to complete.

## III. ACCOUNTING FOR NEW MAP INFORMATION: ON-LINE REPLANNING

In the previous section we studied the issues of planning paths from partial terrain information and the inherent uncertainty in the expected localisation performance along a path. The proposed risk-adverse planning strategy appropriately scales the conservativeness of the planned trajectory to the uncertainty of the map before flight, providing on average the most well-reasoned initial path. It was seen in the previous section that this path can potentially be too conservative in certain scenarios (the planner is risk-adverse, and thus considers worst-case scenarios). As the vehicle actually begins to fly the trajectory and more information about the environment is gathered, aspects of path risk may be reduced based on more certainty of the actual terrain. Additionally, collected information may indicate that the initial planned trajectory is no longer suitable due to a variation between the expected terrain feature density and actual observed terrain feature density; logically the vehicle should update it's initial plan based on new information.

In this section we analyse strategies for determining in-flight when replanning is necessary based on new collected map information and present results of an online replanning scheme.

### A. Detecting When Replanning is Necessary

Once an initial path across the terrain has been planned based on the framework described in [2], the vehicle begins to fly along the trajectory. As the vehicle moves into unmapped areas of the terrain, the on-line replanning system should constantly attempt to assess whether the current plan is expected to meet the localisation performance metric constraints. The planner can predict variations in the original estimated performance by examining two factors: (a) looking for variations in the initial predicted and actual localisation performance metrics or (b) by re-assessing the density of terrain features in newly explored areas, based on the features that are seen. Method (a) provides a more definitive indicator that replanning is necessary with the disadvantage that the localisation performance metric constraints will be violated before replanning is performed. Method (b) provides a less definitive indicator that replanning is necessary but with the advantage that potential violation of constraints can be detected beforehand, with replanning performed in order to prevent violation of the constraints.

### B. Re-Estimating the Density of Terrain Features in Newly Explored Environments

At the initial point of planning, the environment is broken up into regular, 100x100m grids which are labelled as being explored or unexplored based on the presence of initial map features. As the vehicle flies over an unexplored grid, the number of new features observed in the grid ($n_f$) is counted. If the number of features in the grid is outside the range $\frac{1}{2}\rho_{init} < n_f < 2\rho_{init}$ (where $\rho_{init}$ is the initial predicted terrain density) then a potential variation of density in the grid is recorded. If the same density variation is recorded in a number of concurrently explored grids (the number being a tuning parameter of the planner, set equal to three in the following results), then the planner is notified that replanning is required.

### C. Replanning Procedure

The total planning system including both initial planning and on-line adaptive replanning is composed of the following steps:

1) Based on initial terrain information, an initial path to the destination is planned using the framework described in [2].
2) As the vehicle flies over newly explored terrain, the density of features in the new area is computed as described in Section III-B. If the newly computed density lies within the bounds $\frac{1}{2}\rho_{init} < n_f < 2\rho_{init}$ then the vehicle continues to fly along the initial planned path.
3) If a variation in terrain feature density is detected, then a new value for the expected density of features in unexplored areas of the terrain is computed equal to the average of the counted features $n_f$ in the newly explored grids. The localisation metrics are re-predicted along the next trajectory

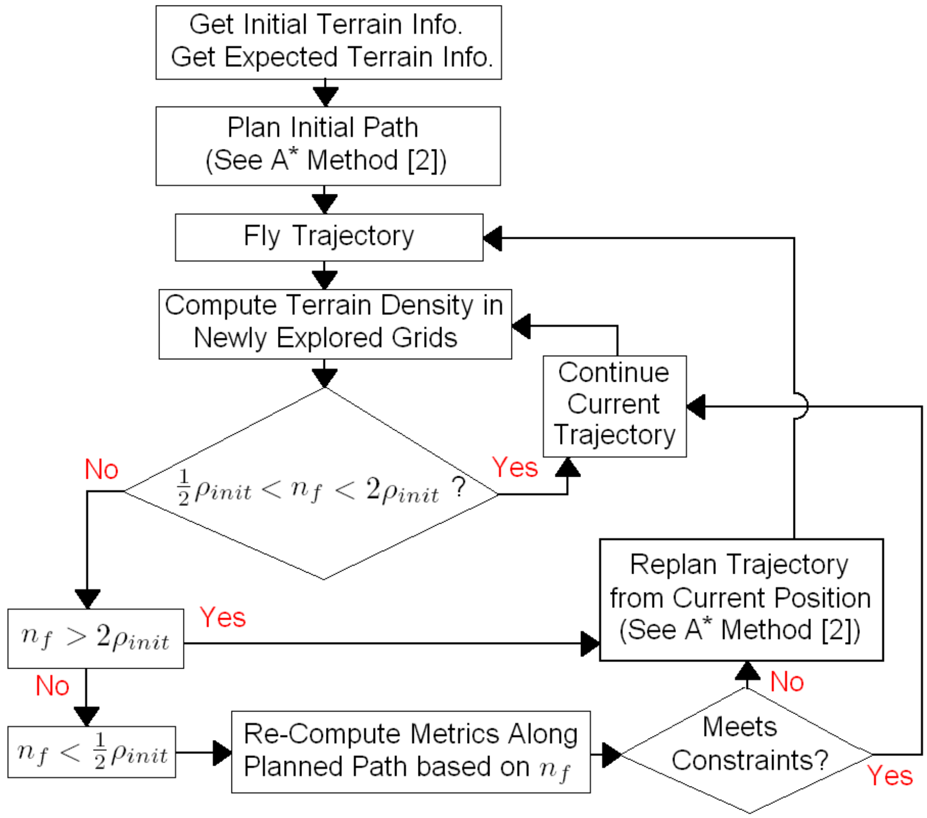


Fig. 10. Online Replanning Process: An initial planned trajectory is formulated and the vehicle begins to fly the trajectory. As the vehicle traverses newly explored terrain, the feature density in new areas is reassessed. If the density is found to rise about a given threshold, a new trajectory to the destination is planned from the current position. If the density is found to fall below a given threshold, the predicted performance along the path is reevaluated based on the new predicted density. If the localisation performance metric constraints are not violated, the vehicle continues on the current planned trajectory. If a violation is predicted, then the trajectory is replanned from the current position, accounting for the lower expected density of features.

segment based on the new value of expected terrain feature density. If a violation of the localisation metric is detected (i.e. in the case where the expected terrain density is now lower), or the expected density of terrain features is now much higher (where it is expected that the current plan may be too conservative) then the trajectory is replanned. Replanning occurs using the method described in [2] from the current vehicle position to the destination based on the new expected terrain feature density in unexplored sections of the terrain.

The replanning process is illustrated in Figure 10.

### *D. Results*

In this section we present results from the online path re-planner when applied in a simulation run over the same initial terrain map as demonstrated in Section II-C. We consider two different scenarios. In both scenarios the structure of the terrain and actual density of features in the unmapped areas in the first half of the map is the same and set to be consistent with the density of known features (set at 2 features per 100x100m area). In the second upper area of the map, the actual density of terrain features was varied; in scenario 1, the density of features was set at 1 feature per 100x100m area, and in scenario 2, the density of features was set at 10 feature per 100x100m area. The online replanning system was tested in both scenarios.

*1) Results Scenario 1: Online Replanning for Lower than Expected Feature Densities:* Figure 11 illustrates the initial and replanned trajectories for scenario 1, where the actual feature densities in the

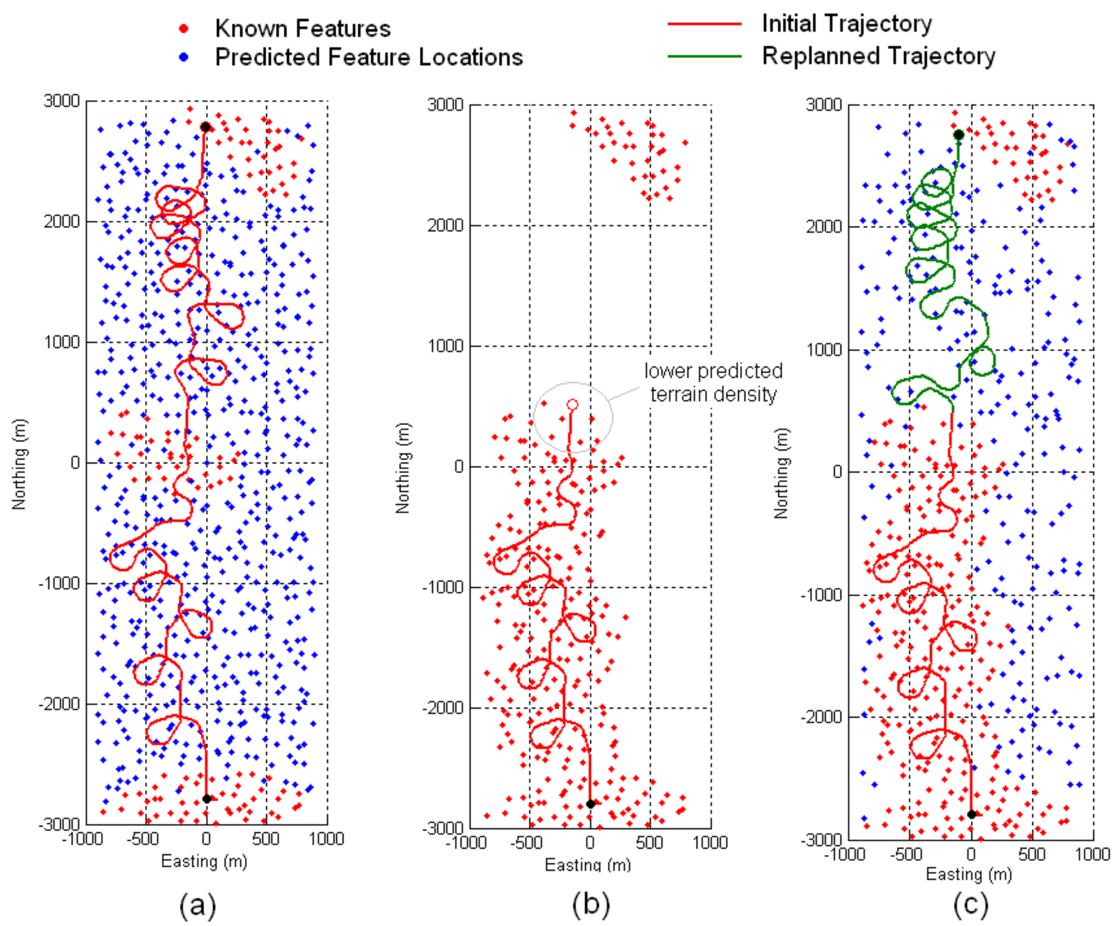


Fig. 11. Initial and Replanned Trajectories for Scenario 1 where Lower than Expected Feature Density is Discovered: (a) The initial planned trajectory over the terrain given an expected feature density of 2 features per 100x100m. Prior known features are shown in red and expected feature locations in blue. (b) The actual flown trajectory and discovered features up to the point before replanning. (c) The replanned trajectory shown in green. The blue features here indicate the re-projected expected features into the remaining unknown locations at a new expected density of 1 feature per 100x100m.

second half of the map are lower than expected. Figure 11 (a) demonstrates the initial planned path at the start of the mission. Known areas of the map are illustrated by the red points whereas expected map features in unmapped areas are illustrated by the blue points. The vehicle begins to fly the trajectory and at a point about 200 seconds into the flight, the measured terrain density of the newly explored regions drops below the threshold of $\frac{1}{2}\rho_{init}$ =1 feature per 100x100m area (actual path and feature illustrated in Figure 11 (b)). At this point the localisation metrics are predicted forward for the remaining sections of the trajectory based on a new expected terrain density of 1 feature per 100x100m in the remaining unmapped areas. A violation of the constraints is detected and the thus the trajectory is replanned to the destination. Figure 11 (c) illustrates the replanned segment of the trajectory with the newly computed expected positions of features in the unmapped areas (based on a density of 1 feature per 100x100m) shown by the blue points. Note that the path length of the new trajectory is now longer with more re-visiting or looping behaviors (more conservative) owing the the reduced density of expected features. The total path now takes 425 seconds to complete as opposed to 415 seconds, as predicted at the beginning of the mission.

Figure 12 illustrates both the expected values for the metrics at the beginning of the trajectory (where the predicted density of features in unmapped areas is 2 features per 100x100m) (shown in blue), and the updated prediction of metrics along the pre-planned path based on the updated expected feature density of 1 feature per 100x100m (shown in red). The performance metrics for the initial planned path (predicted

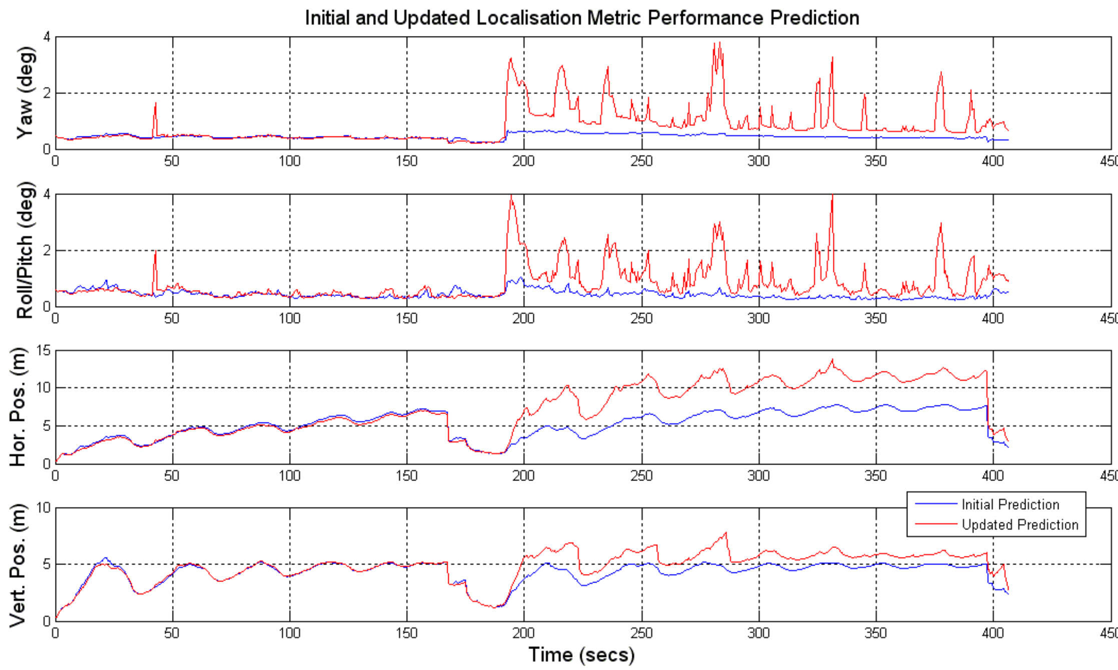


Fig. 12. Initial Prediction and Updated Prediction of Localisation Metric Performance for Scenario 1 where Lower than Expected Feature Density is Discovered: The blue curve indicates the expected values for the metrics at the beginning of the trajectory (where the predicted density of features in unmapped areas is 2 features per 100x100m). The red curve indicates the updated prediction of metrics along the pre-planned path based on the updated expected feature density of 1 feature per 100x100m. The updated prediction indicates that the pre-planned path is now expected to violate the localisation constraints and thus the path should be replanned.

at the start of the mission) lie below the constraints, and thus at the start of the mission, this trajectory is considered adequate. When the vehicle reaches a point about halfway through the map, the measured density of terrain features in newly explored areas drops from 2 features per 100x100m to 1 feature per 100x100m. The re-predicted localisation performance metrics now violate the localisation constraints, and the path is replanned from the current point.

Figure 13 illustrates the updated localisation performance metrics for the replanned trajectory from the point of replanning. The metric values are now decreased for the more conservative, replanned trajectory where the localisation error constraints are no longer violated.

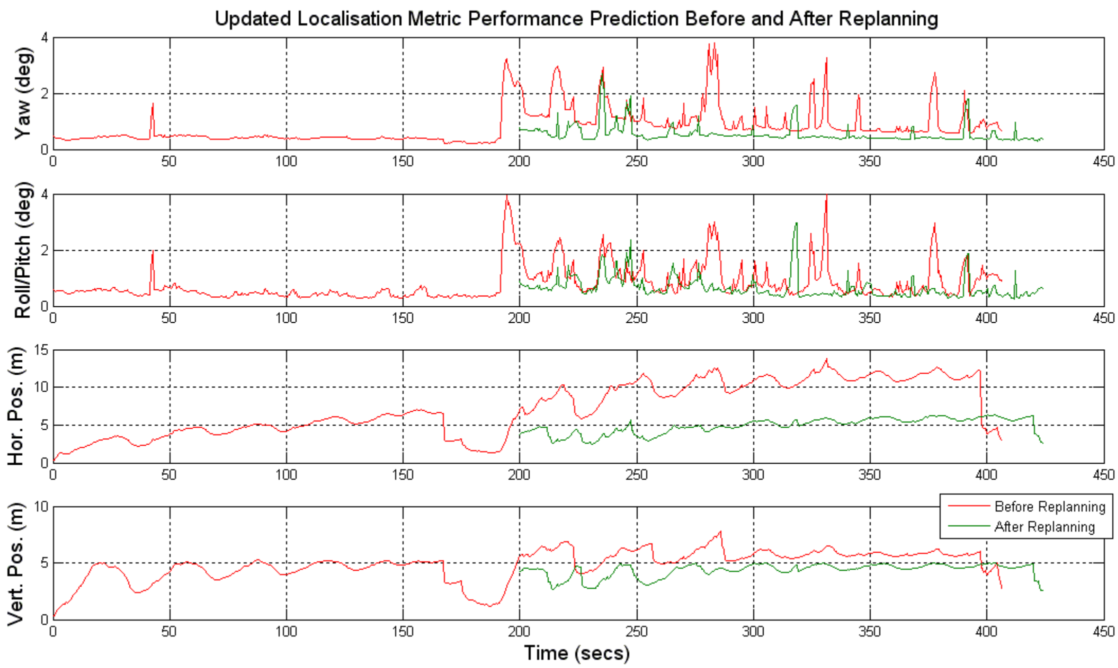


Fig. 13. Updated Prediction Before and After Trajectory Replanning for Scenario 1 where Lower than Expected Feature Density is Discovered: The red curve indicates the updated prediction of metrics along the pre-planned path. The green curve indicates the update of the predicted metrics after the path has been replanned. The predicted performance metrics for the replanned path no longer violate the localisation constraints, and the vehicle can proceed along the path.

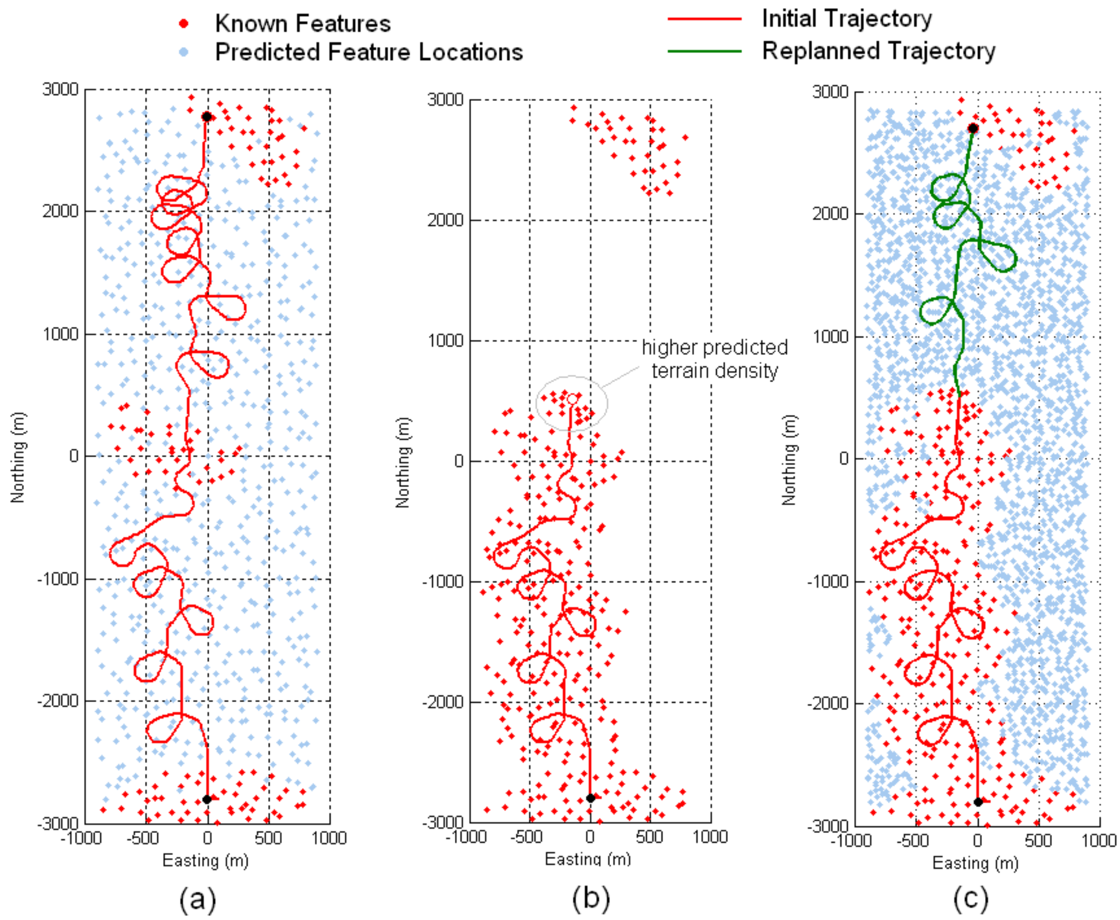


Fig. 14. Initial and Replanned Trajectories for Scenario 2 where Higher than Expected Feature Density is Discovered: (a) The initial planned trajectory over the terrain given an expected feature density of 2 features per 100x100m. Prior known features are shown in red and expected feature locations in blue. (b) The actual flown trajectory and discovered features up to the point before replanning. (c) The replanned trajectory shown in green. The blue features here indicate the re-projected expected features into the remaining unknown locations at a new expected density of 10 features per 100x100m.

*2) Results Scenario 2: Online Replanning for Higher than Expected Feature Densities:* Figure 14 illustrates the initial and replanned trajectories for scenario 2, where the actual feature densities in the second half of the map are higher than expected. As in Scenario 1, Figures 14 (a) and (b) illustrate the pre-planned path and the path up to the point of replanning respectively. At this point in the trajectory (around the 200 second mark), the measured terrain density of the newly explored regions rises above the threshold of $2\rho_{init}$ =4 features per 100x100m area to a measured value of 10 features per 100x100m area, and thus the trajectory is replanned from this point. Figure 14 (c) illustrates the replanned segment of the trajectory with the newly computed expected positions of features in the unmapped areas (based on a density of 10 features per 100x100m) shown by the blue points. Note that the path length of the new trajectory is now shorter and the vehicle reaches the destination in less time owing to the ability to 'cut-out' some of the looping and revisiting behaviors due to the increase in expected terrain feature density.

Figure 15 illustrates both the expected values for the metrics at the beginning of the trajectory (where the predicted density of features in unmapped areas is 2 features per 100x100m) (shown in blue), and the updated prediction of metrics along the pre-planned path based on the updated expected feature density of 10 features per 100x100m (shown in red). The updated prediction of localisation performance metrics indicate that the original planned trajectory over the terrain will meet the localisation constraints but that expected accuracy is much larger than is necessary (expected errors are smaller).

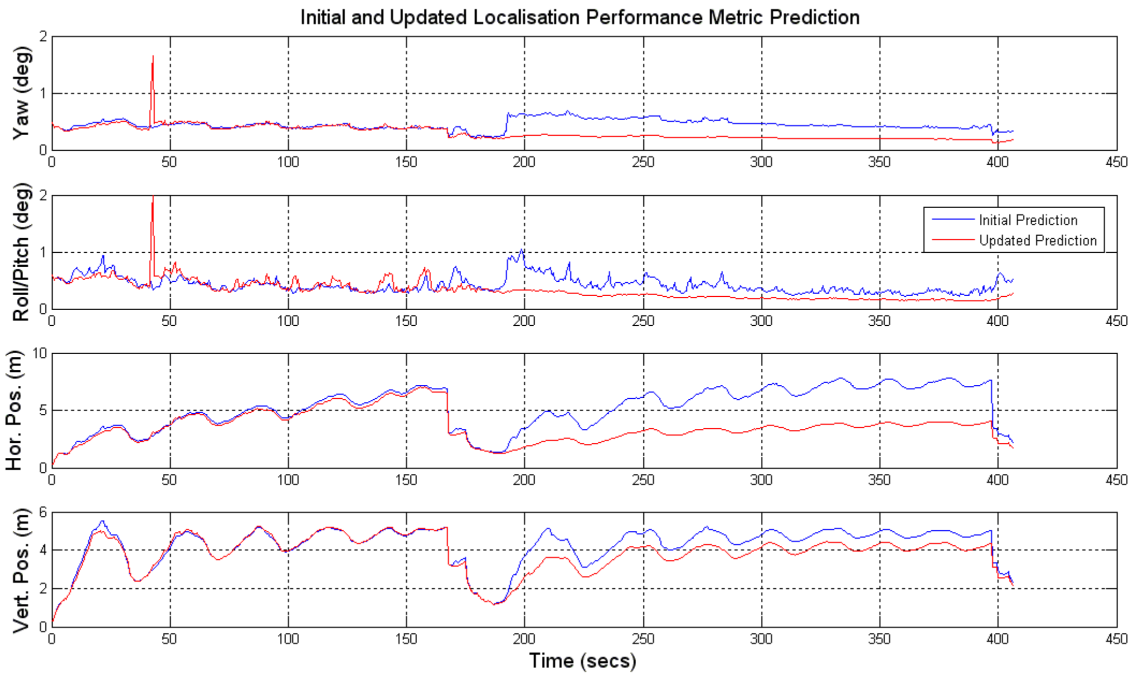


Fig. 15. Initial Prediction and Updated Prediction of Localisation Metric Performance for Scenario 2 where Higher than Expected Feature Density is Discovered: The blue curve indicates the expected values for the metrics at the beginning of the trajectory (where the predicted density of features in unmapped areas is 2 features per 100x100m). The red curve indicates the updated prediction of metrics along the pre-planned path based on the updated expected feature density of 10 features per 100x100m. The updated prediction indicates that the pre-planned path is now too conservative and that potentially a more direct trajectory to the destination point that does not violate the constraints exists.

Figure 16 illustrates the updated localisation performance metrics for the replanned trajectory from the point of replanning. The metric values are maintained below the constraints for the newly planned path and the length of the path has now been significantly reduced; the vehicle is expected to arrive at the destination point 340 seconds after the beginning of the mission as opposed to 410 seconds after the beginning of the mission if the vehicle had not replanned the trajectory.

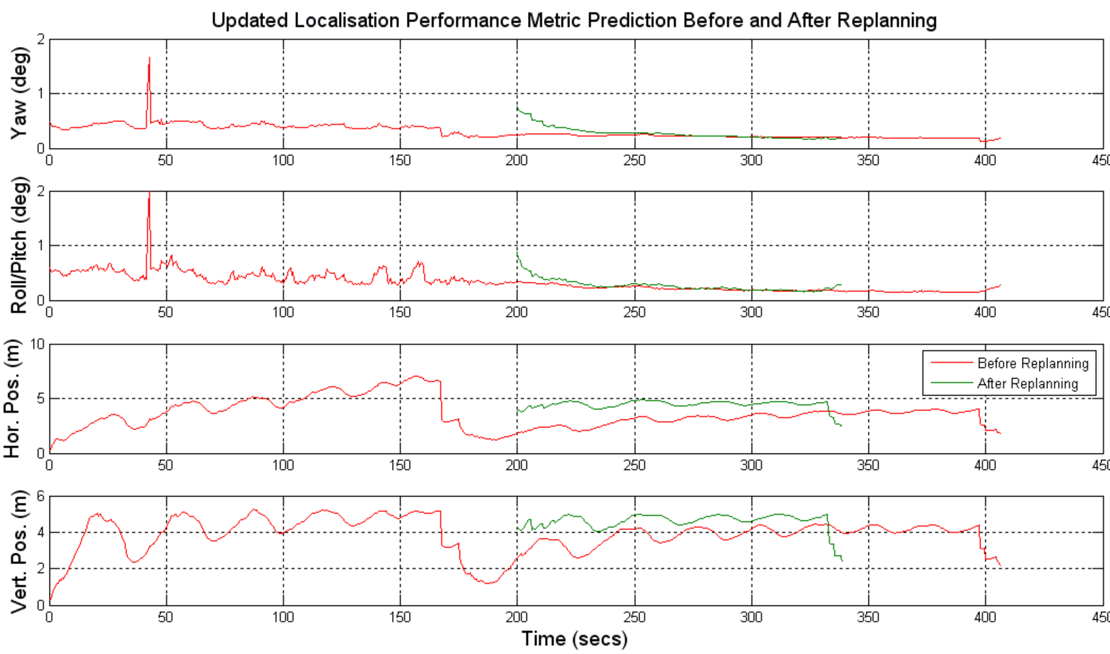


Fig. 16. Updated Prediction Before and After Trajectory Replanning for Scenario 2 where Higher than Expected Feature Density is Discovered: The red curve indicates the updated prediction of metrics along the pre-planned path. The green curve indicates the update of the predicted metrics after the path has been replanned. The predicted performance metrics for the replanned path remain within the localisation error constraints where the planned path now arrives at the destination point in less time than the original trajectory.

## IV. Discussion, Conclusions and Future Work

In this section we summarise and provide a discussion for the results demonstrated in the previous sections.

### A. Accounting for Uncertainty in Localisation Performance over Unmapped Terrain

In the first half of this project [2], it was recognised that terrain-aided localisation system performance while flying over unmapped areas of the terrain required an approximation of the density of terrain features in this area and thus resulted in a degree of uncertainty of the performance. It was proposed that the magnitude of this uncertainty and it's repercussions to the planning process should be studied, along with methods for accounting for the uncertainty or risk in paths planned over unmapped terrain.

In the first section of the work, we considered the ramifications to the path planning process of localisation performance metric uncertainty during prediction over unmapped terrain. Monte carlo simulation results in Section II-B indicated the potential uncertainty for different types of trajectories and for different values of expected terrain density in the un-mapped areas. It was found that performance metrics predicted over high-density terrain were reasonably repeatable and invariant to the exact configuration of features, whereas metrics predicted over low-density terrain were highly variant and sensitive to the exact configuration of features. In order to mitigate the uncertainty or risk involved in making trajectory plans based on this information, a risk-adverse extension to the path planning system was proposed in Section II-C which varied performance metric constraints to scale to the level of uncertainty in the predicted performance metric values.

Future work should look towards more accurate methods for predicting localisation performance variance based on aspects of the mission (other than feature density) such as UAV and sensor configurations (i.e. accuracy of the IMU and terrain sensors). Potential avenues for this type of prediction would probably lie in a reformulation of the predicted state space away from having to predict individual feature contributions, towards modelling localisation covariance as an explicit function of feature density.

### B. Accounting for New Map Information: On-line Replanning

In the first half of the project [2], the proposed planning methods were limited to initial path planning at the start of a UAV mission. It was proposed that the methods should be extended to provide the ability to replan and alter paths on-line, while the vehicle is in-flight based on new terrain information that becomes available, providing a robust and adaptable planner.

In the second section of the work, we considered methods for online, in-flight replanning of pre-planned trajectories. Based on the knowledge of the link between localisation performance in SLAM and the actual density of features in the terrain, it was proposed to use detected variations in the density of the terrain as an indicator for the need to re-plan trajectories. The resulting replanned trajectories were shown to meet localisation constraints and decrease the time taken to arrive at the destination when the density of terrain features in unmapped areas varied from one location to another.

Future work into on-line replanning should be directed in two areas. Firstly, other more accurate indicators of performance variation (other than detected change in feature density) should be investigated to provide more information on when and how to replan during flight. Secondly, more efficient methods should be investigated for making small variations to planned paths, avoiding the need to replan all of the remaining segments. This is difficult to achieve when localisation performance in SLAM is used as a constraint to the planner as small changes mid-way through the flight can potentially have performance effects further into the flight based on skipping observations of features or missing loop closure opportunities. Potential avenues for solving this problem may include the use of a $D^*$ search algorithm [5] with adaptation to the active-SLAM framework and localisation performance constraints.

### C. Other Avenues for Future Work in Cooperative SLAM

The multi-vehicle version of the path planning algorithms presented in this project have so far only been developed for a centralised path planning system. Work on developing a decentralised or distributed version of these algorithms is a logical next step in the research. There are several factors that complicate the development of distributed versions of planning algorithms based on an $A^*$ or $D^*$ search related to the tight-coupling of integral constraints in the localisation constraints-based path planner presented in this work. Planning in a decentralised manner means that when constraints are determined to be violated in one vehicle's path, potentially each other vehicle in the team must reevaluate their trajectory plans associated to this vehicle; the computational/communications complexity of this type of architecture would be prohibitive. Decentralisation of the algorithm would require a shift in the way performance metrics were calculated, moving away from the current idea of using covariance matrix propagation.

The path planning algorithms developed in this project rely on the evaluation of SLAM joint covariance matricies that represent the expected levels of uncertainty (and thus accuracy) in different estimated states. Unfortunately, the computational complexity of covariance matrix evaluation in SLAM is known to be $\mathcal{O}(N^2)$ with the size of the map. The computational complexity of the current approach, although suitable for planning over small-to-medium distances and with a handful of vehicles, will become very large as the size of the planning distance and the number of vehicles increases. One potential avenue for developing planning algorithms that scale well to larger distances and numbers of vehicles is through the use of a sparse information form representation of SLAM joint estimates, where potentially the computational complexity of creating plans can be reduced to $\mathcal{O}(N)$, linear to the size of the planning distance and number of vehicles.

## REFERENCES


[1] M. Bryson and S. Sukkarieh. Final Report - Cooperative Airborne Inertial-SLAM for Improved Platform and Feature/Target Localisation. In *USAF/AOARD contract FA5209-05-P-0550*, 2007.

[2] M. Bryson and S. Sukkarieh. Mid-term Report - Extending Cooperative SLAM into Multi-Objective Missions. In *USAF/AOARD contract FA4869-08-1-4060*, 2008.

[3] M. Bryson and S. Sukkarieh. Observability Analysis and Active Control for Airborne SLAM. *IEEE Transactions on Aerospace and Electronic Systems*, 44(1):261–280, 2008.

[4] R. Dechter and J. Pearl. Generalized best-first search strategies and the optimality of A*. *Journal of the Association of Computer Machinery*, 32(3):505–536, 1985.

[5] A. Stentz. Optimal and efficient path planning for partially-knownenvironments. In *IEEE International Conference on Robotics and Automation*, 1994.

[6] S. Sukkarieh, M. Bryson, and R. Hardwick-Jones. Mid-term Report - Cooperative Airborne Inertial-SLAM for Improved Platform and Feature/Target Localisation. In *USAF/AOARD contract FA5209-05-P-0550*, 2006.